arXiv:1412.3721v1 [cs.DS] 11 Dec 2014

Approximation Algorithms for Budget Constrained Network Upgradeable Problems

Debjyoti Saharoy and Sandeep Sen

Department of Computer Science and Engineering,
IIT Delhi, New Delhi 110016, India*

December 12, 2014

Abstract

We study *budget constrained network upgradeable problems.* We are given an undirected edge weighted graph $G = (V, E)$ where the weight an edge $e \in E$ can be upgraded for a cost $c(e)$. Given a budget B for improvement, the goal is to find a subset of edges to be upgraded so that the resulting network is optimum for B. The results obtained in this paper include the following.

1. *Maximum Weight Constrained Spanning Tree*

   We present a randomized algorithm for the problem of weight upgradeable budget constrained maximum spanning tree on a general graph. This returns a spanning tree $\mathcal{T}'$ which is feasible within the budget B, such that $\mathbf{Pr}[l(\mathcal{T}') \geq (1-\epsilon)\text{OPT}\ ,\ c(\mathcal{T}') \leq B] \geq 1 - \frac{1}{n}$ (where l and c denote the length and cost of the tree respectively), for any fixed $\epsilon > 0$, in time polynomial in $|V| = n$, $|E| = m$. Our results extend to the minimization version also. Previously Krumke et. al. [KNM$^+$97] presented a $(1 + \frac{1}{\gamma}, 1 + \gamma)$ bicriteria approximation algorithm for any fixed $\gamma > 0$ for this problem in general graphs for a more general cost upgrade function. The result in this paper improves their 0/1 cost upgrade model.

2. *Longest Path in a DAG*

   We consider the problem of weight improvable longest path in a n vertex DAG and give a $O(n^3)$ algorithm for the problem when there is a bound on the number of improvements allowed. We also give a $(1-\epsilon)$-approximation which runs in $O(\frac{n^4}{\epsilon})$ time for the budget constrained version. Similar results can be achieved also for the problem of shortest paths in a DAG.

1 Introduction

We consider optimization problems in budget constrained network *upgradeable* model, where the exact cost-benefit trade-offs are known[1] a priori. In such optimization problems, there are two kinds of problem data: First, the nominal or unimproved data, and second, *improved* data of several stages that would allow solutions with better performance. We are allowed to use the improved data by incurring a cost within an overall budget for all the improvement cost. Several computational problems fall in this category. For example, consider the problem of upgrading arcs in a network to minimize travel time [CLZ06]. Another example would be a variant of min cost flow problem.

*Email addresses: {debjyoti,ssen}@cse.iitd.ac.in

[1] which may not hold for an online model

Here we are allowed to lower the flow cost of each arc, and given a flow value and a bound on the total budget which can be used for lowering the flow costs, the goal is to find an upgrade strategy and a flow of minimum cost. This problem is considered by [DNW02]. In [KNW$^+$99] Krumke et. al. consider the problem of improving spanning trees by upgrading nodes. In communication networks upgrading could mean installing faster communication device at a node. Typically, such an upgradeable extension increases the complexity of an otherwise polynomially solvable problem to being NP-hard. The constrained minimum spanning tree problem subject to an additional linear constraint was shown to *weakly* NP Hard in [AAN82].

Remark In the remaining paper, we will use the term *improvable* instead of *upgradeable* as a common term to address both increase or decrease in the parameter value, (as the case may be), to obtain a superior or improved objective value.

The *improvable* version of a problem is distinct from optimization version that may have limited choices for certain kinds of objects. For example, we may have multiple (parallel) edges for the network problems offering trade-offs between cost and weight. This can be captured within the framework of the original problem, using additional constraints, and keeping the objective function unchanged. The improvable version works on a budget constraint that is distinct from the original constraints and there may not be any well-defined conversion between the budget constraint and the original constraints. For example, in the constrained spanning tree problem in [AAN82], the lengths and the weights are distinct parameters.

Computing a minimum or maximum weight spanning tree for a network is a well studied problem in computer science. Here we consider a variant of the maximization version of the problem, where the goal is to increase the edge lengths of a given network so that the length of a maximum spanning tree in resulting network is as large as possible. The problem is considered in a context where there is a cost associated with improving an edge and there is a budget constraint on the total cost of improving the edges. We denote the problem as $IMST$ for *Improvable Maximum Spanning Tree*, the optimum tree as $\mathcal{T}^*$, its length as $l(\mathcal{T}^*)$ and the total cost of the higher length valued edges as $c(\mathcal{T}^*)$. A precise definition of the problem is given in section 2. Since the maximum weight base problem on a matroid is equivalent to minimum cost base problem (by applying suitable transformations on weights) so our results and analysis hold even for weight improvable budget constrained minimum spanning tree on general graphs.

The second problem that we consider is the budget constrained network improvable longest or shortest path in a directed acyclic graph. Since the longest and shortest path problems in a DAG are related, we consider the former but similar results and analysis would hold for the later as well. The longest path problem is the problem of finding a simple path of maximum length in a graph. The longest path problem is NP-Hard and there is no polynomial time constant-factor approximation algorithm unless P=NP ([KMR97]). However, the problem becomes polynomial time solvable on directed acyclic graphs. We investigate this problem in budget constrained network improvable setting. We refer to the problem as $WILDAG$ for *Weight-Improvable-Longest-Path-in-DAG*. Longest path algorithms have applications in diverse fields. The well known Travelling Salesman Problem is a special case of the Longest Path Problem ([HN62]). The longest path in a program activity graph is known as the critical path which represents the sequence of program activities that take the longest time to execute. Longest path algorithms are required to calculate critical paths.

1.1 Other Related Work

More examples and applications of computational problems in the improvable framework can be found in [GSSS14]. Goerigk, Sabharwal, Schöbel and Sen [GSSS14] considered the weight-reducible knapsack problem, for which they gave a polynomial-time 3-approximation and an FPTAS for the

special case of uniform improvement costs. The problem of budget constrained network improvable spanning tree has been proved to be NP-hard, even for series-parallel graphs, by Krumke et. al. [KNM+97], which also cite several practical applications. Frederickson and Solis-Oba [FSO96] considered the problem of increasing the weight of the minimum spanning tree in a graph subject to a budget constraint where the cost functions increase linearly with weights. Berman et. al. [BIO92] consider the problem of shortening edges in a given tree to minimize its shortest path tree weight. In contrast to most problems in the network upgradation model, this problem was shown to be solvable in strongly polynomial time. Phillips [Phi93] studied the problem of finding an optimal strategy for reducing the capacity of the network so that the residual capacity in the modified network is minimized.

Most of the network modification results can be broadly classified as bi-criteria problems which are characterized as (α, β) approximation if the algorithm achieves factor α (respective β) approximation w.r.t. to the first (second) parameter. If α (or β) equals 1, then it is at least as good as the optimal solution w.r.t. to the first (respectively second) parameter. For the improvable spanning tree problem, the two parameters are the total weight of the spanning tree and the budget available for improving the spanning tree edges. Krumke et al. [KNM+97] show that Tree-width bounded graphs with linear reduction costs are $(1+\epsilon, 1+\xi)$-approximable for any fixed ϵ, $\xi > 0$. They also show general graphs are $(1+\frac{1}{\gamma}, 1+\gamma)$-approximable for any fixed $\gamma > 0$ which implies a trade-off between the two approximation factors (note that the balance occurs for $\gamma = 1$). In the more elaborate journal version, the authors [DKNMR98] actually consider three distinct models of improvements where the 0/1 reduction comes closest to this paper. Ravi and Goemens [RG96] have studied the constrained minimum spanning tree problem with two independent weight function on the edges and gave a $(1, 1+\epsilon)$ polynomial time approximation. Their method can be used to derive a $(1+\epsilon, 1)$ approximation for the constrained minimum spanning tree problem that runs in pseudo-polynomial time ([MRS+98]).

1.2 Our Contributions

For the problem of *IMST* which allows multistage improvements, we present a randomized algorithm that returns a spanning tree $\mathcal{T}'$ which is feasible within the budget B and has length at least $(1-\epsilon)$ times the OPT[2] with high probability, i.e $\mathbf{Pr}[l(\mathcal{T}') \geq (1-\epsilon)\text{OPT}\ ,\ c(\mathcal{T}') \leq B] \geq 1 - \frac{1}{n}$, for any fixed $\epsilon > 0$, in time polynomial in $|V| = n$, $|E| = m$. Our algorithm does not make any assumptions on the structure of the graph and works for general graphs.

For the problem of *WILDAG* we give an $O(n^3)$ algorithm for the special case of uniform improvement cost for each edge. We also consider the more general version with arbitrary improvement costs and a budget constraint on the total improvements and give a fully polynomial time approximation scheme for this problem.

The primary observation that we exploit is that the optimal solution comprises of the improved and unimproved versions in some ratio which is not known but can be approximated in some way. For example, if we restrict ourselves to the unweighted version (of the maximization problem), we can say that at least half of the objective value is due to the improved or the unimproved elements. If we know how to solve each version separately, even approximately, we can combine them to obtain a constant factor approximation overall. This idea extends to the weighted version also. We also make use of dynamic programming and scaling to obtain our results. Our techniques are quite general and likely to be useful for other improvable problems.

[2]The value of the optimum spanning tree among those trees that use up a total improvement budget B

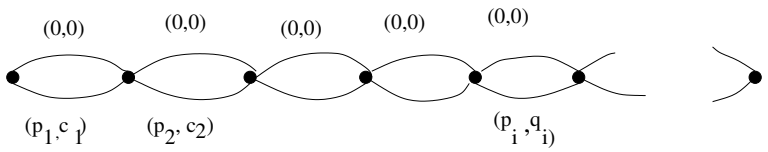

Figure 1: Reducing Knapsack to IMST and WILDAG

1.3 Organization of the paper

In section 2 we investigate the problem $IMST$. Beginning with a simple $\frac{1}{2}$-approximation for the special case of uniform improvement costs in section 2.1, we consider more general improvement costs in section 2.2 and give a randomized bi-criteria $(1-\epsilon, 1)$-approximation. We extend the algorithm in 2.2 to handle multistage improvements in 2.3.

In section 3, we consider the problem of longest path in DAG (3) in the improvable framework. For a gentle exposition of the ideas and techniques, we first deal with the $WILDAG$ problem for the special case of equal improvement cost for each edge and give an $O(n^3)$ algorithm. In section 3.1 we consider the general version with non-uniform improvement costs and a budget constraint and present a fully polynomial time approximation scheme in section 3.2.

2 Randomized $(1-\epsilon)$-approximation for $IMST$

Given an undirected graph $G = (V, E)$, and nonnegative integers l_e and h_e (i.e each edge can be thought to have two copies with length l_e and h_e), an improvement cost c_e, for each edge $e \in E$,and a budget B, we consider the problem of finding a spanning tree with maximum total edge length under the restriction that each time a higher valued edge is used, the associated improvement cost is incurred and the total improvement cost can be at most B. Similar to the constrained MST problem, this variation is also computationally intractable. For completeness, we present a simple proof along the lines of [DKNMR98, AAN82].

Theorem 1 *The Improvable MST (*IMST*) problem is NP hard even when the graph is a tree.*

Proof: We can reduce the knapsack problem using a construction similar to [AAN82] - see Figure 1. The knapsack profit function is $\max \sum_e p_e x_e$ and the constraint is $\sum_e c_e \cdot x_e \leq B$ where B is also the budget constraint of the IMST problem. Wlog, we have shown the two different versions of the edges as parallel edges, since any spanning tree will include exactly one of the two edges. Note that $l_e = 0$ and $h_e = p_e$. The IMST solution will directly yield a solution to the knapsack problem. □

2.1 Uniform improvement

If the improvement costs are uniform, i.e., $c_e = 1$, then the problem reduces to finding a spanning tree with maximum length subject to a cardinality constraint (say k) on the number of higher length valued edge used. Let us denote the problem as $UIMST$ for *Uniform Improvable Maximum Spanning Tree*, the optimum tree as $\mathcal{T}^*$, and its length as $l(\mathcal{T}^*)$.

2.1.1 The Algorithm

Note that the optimum tree is composed of two spanning forests one with edges having lengths l_e and the other with lengths h_e (at most k of them). We denote the two forests as $\mathcal{F}_1^*$ and $\mathcal{F}_2^*$ and

their respective lengths as $l(\mathcal{F}_1^*)$ and $l(\mathcal{F}_2^*)$. Our algorithm is based on finding two spanning trees $\mathcal{T}_1$ and $\mathcal{T}_2$ such that

$$l(\mathcal{T}_1) + l(\mathcal{T}_2) \geq l(\mathcal{F}_1^*) + l(\mathcal{F}_2^*) = l(\mathcal{T}^*)$$

So that $\max\{l(\mathcal{T}_1), l(\mathcal{T}_2)\} \geq \frac{1}{2} l(\mathcal{T}^*)$. Thus the better of the two trees $\mathcal{T}_1$, $\mathcal{T}_2$ is a $\frac{1}{2}$-approximate solution of *UIMST*. We can obtain $\mathcal{T}_1$ by invoking Kruskal's algorithm for maximum spanning tree [Kru56] on the lower length valued edges, i.e l_e for each $e \in E$. Clearly $l(\mathcal{T}_1) \geq l(\mathcal{F}_1^*)$. To find the tree $\mathcal{T}_2$ we first invoke Kruskal's algorithm [Kru56] on the high valued edges, i.e h_e for each $h_e \in E$, to obtain the maximum spanning forest in G with at most k edges where k is the maximum number of improvements allowed (we terminate the greedy algorithm after k edges have been chosen). This forest can be extended to a tree $\mathcal{T}_2$ as spanning trees form the bases of a Graphic Matroid. Again, we have $l(\mathcal{T}_2) \geq l(\mathcal{F}_2^*)$. This completes the algorithm of section 2.1.1.

2.2 Arbitrary improvement: An approximation algorithm

We now address the general case where c_e's can be arbitrary. Let us denote the the maximal tree by $\mathcal{T}^*$, its length as $l(\mathcal{T}^*)$ and the total cost of the higher length valued edges as $c(\mathcal{T}^*)$.

Before we formally present the algorithm, we will need the following definition and a related result.

Definition 2 *Two Cost Spanning Tree Problem: Given a connected undirected graph $G = (V, E)$, two edge weight functions, $c()$ and $l()$, and a bound B, find a spanning tree $\mathcal{T}^*$ of G such that the total cost $c(\mathcal{T}^*)$ is at most B and the total cost $l(\mathcal{T}^*)$ is maximum among all spanning trees that stay within the budget constraint.*

Observe that the definition is not completely symmetric in the two weight functions; only the budget B is specified. This problem has been addressed earlier and the following bi-criteria approximation result was presented.

Theorem 3 ([RG96]) *For all $\epsilon \geq 0$, there is a polynomial time approximation algorithm for the Two Cost Spanning Tree problem with a performance of $(1, 1+\epsilon)$.*

The result of Theorem 3 holds even if the set of spanning trees are replaced by the bases of any matroid. The reader may note that since the definition of the Two Cost Spanning Tree is not symmetric in the two weight functions, the $(1, 1+\epsilon)$ approximation does not imply a $(1+\epsilon, 1)$ approximation for the constrained minimum spanning tree problem. Indeed the the running time becomes pseudo polynomial ([MRS$^+$98]).

Let $G' = (V, E')$ be a multi-graph obtained from G, where E' is the set of edges that two copies of an edge $e \in G$ with weights h_e and l_e respectively. The edges h_e have cost c_e and edges l_e have costs 0, for each $e \in E'$. It is easy to see that multi-graph G' retains the properties of a graphic matroid. So we run the algorithm by [RG96] on the multi-graph G' with the budget constraint B. The output of [RG96] is a tree $\mathcal{T}$ such that $l(\mathcal{T}) \geq l(\mathcal{T}^*)$. Note that $\mathcal{T}$ is composed of of a forest of high valued edges h_e, $\mathcal{F}_1$, and a forest of low valued edges l_e, $\mathcal{F}_2$. We must have $c(\mathcal{F}_1) \leq (1+\epsilon)B$ and $c(\mathcal{F}_2) = 0$. To make the forest $\mathcal{F}_1$ feasible within the budget B, we sample the edges in $\mathcal{F}_1$ randomly with probability $\frac{1}{(1+\epsilon)^2}$. That is, with probability, $\frac{1}{(1+\epsilon)^2}$ we retain the high valued edge $h_e \in \mathcal{F}_1$ otherwise we pick the low valued copy of the edge i.e l_e. We denote this randomly sampled subset of edges (which is a forest) $\mathcal{F}_1'$. As $\mathbb{E}(l(\mathcal{F}_1')) = \frac{1}{(1+\epsilon)^2} l(\mathcal{F}_1)$ and $\mathbb{E}(c(\mathcal{F}_1')) \leq \frac{1}{(1+\epsilon)} B$, we take the union of the two forests $\mathcal{F}_1'$ and $\mathcal{F}_2'$ that forms a tree $\mathcal{T}'$, since the edges are in 1-1 correspondence

with $\mathcal{T}$. Then

$$\begin{aligned}\mathbb{E}(l(\mathcal{T}^{'})) &= \mathbb{E}(l(\mathcal{F}_1^{'})) + l(\mathcal{F}_2) \\ &= \frac{1}{(1+\epsilon)^2} l(\mathcal{F}_1) + l(\mathcal{F}_2) \\ &\geq \frac{1}{(1+\epsilon)^2}(l(\mathcal{F}_1) + l(\mathcal{F}_2)) = \frac{1}{(1+\epsilon)^2} l(\mathcal{T})\end{aligned}$$

Hence, we have

$$\mathbb{E}(l(\mathcal{T}^{'})) \geq \frac{1}{(1+\epsilon)^2} l(\mathcal{T}^*) \tag{1}$$

$$\mathbb{E}(c(\mathcal{T}^{'})) \leq \frac{1}{(1+\epsilon)} B \tag{2}$$

We have shown that our algorithm returns a spanning tree whose expected cost is feasible and has expected length at least $\frac{1}{(1+\epsilon)^2}$ times the optimum. We shall claim the following result.

Theorem 4 *Given an undirected graph $G = (V, E)$, an error parameter ϵ and a confidence parameter δ, the algorithm of section 2.2 returns a spanning tree $\mathcal{T}^{'}$ which is feasible within the budget B, such that*

$$\mathbf{Pr}[l(\mathcal{T}^{'}) \geq (1-\epsilon) l(\mathcal{T}^*) \ , \ c(\mathcal{T}^{'}) \leq B] \geq 1 - \delta,$$

in time polynomial in $|V| = n$, $|E| = m$, $\frac{1}{\epsilon}$ and $\log(\frac{1}{\delta})$.

Proof: It suffices to achieve the above with $\delta = c$, where $c < 1$ is any arbitrary constant. The error probability can be later boosted to the given δ by performing $O(\log \frac{1}{\delta})$ trials with error probability c and taking the median.

Now to run the algorithm in section 2.2 let us choose ϵ' (note that ϵ' is the parameter of the algorithm in [RG96] and ϵ is the desired error guarantee) such that $1 - \epsilon = (1 - \epsilon')^2$. Therefore,

$$\begin{aligned}& l(\mathcal{T}^{'}) \leq (1-\epsilon) l(\mathcal{T}^*) \\ & \Leftrightarrow l(\mathcal{T}^{'}) \leq (1-\epsilon')^2 \mu (1+\epsilon')^2 (\text{where } \mu = \mathbb{E}(l(\mathcal{T}^{'}))) \\ & \Leftrightarrow l(\mathcal{T}^{'}) \leq (1-\epsilon'^2)^2 \mu \leq (1-\epsilon'^2)\mu\end{aligned}$$

The first implication follows from Equation 1 in section 2.2. Therefore,

$$\begin{aligned}\mathbf{Pr}[l(\mathcal{T}^{'}) \leq (1-\epsilon) l(\mathcal{T}^*)] &\leq \mathbf{Pr}[l(\mathcal{T}^{'}) \leq (1-\epsilon'^2)\mu] \\ &\leq \mathbf{exp}(-\frac{\epsilon'^4 \mu}{3}) \leq \mathbf{exp}(-\frac{\epsilon'^4 \frac{1}{(1+\epsilon')^2} l(\mathcal{T}^*)}{3})\end{aligned}$$

The second inequality follows from Chernoff bounds and the third inequality from Equation 1. If $\mathbf{exp}(-\frac{\epsilon'^4 \frac{1}{(1+\epsilon')^2} l(\mathcal{T}^*)}{3}) < \mathbf{exp}(-1)$ then $\mathbf{Pr}[l(\mathcal{T}^{'}) \leq (1-\epsilon) l(\mathcal{T}^*)] < \frac{1}{e}$. If $l(\mathcal{T}^*) > \frac{3(1+\epsilon')^2}{\epsilon'^4} = s$, we are done, otherwise, we can scale the edge lengths according to the following technical Lemma.

Lemma 5 *The edge lengths of the graph can be additively scaled by $\frac{s}{n}$ while preserving the $(1-\epsilon)$ approximation guarantee.*

Proof: From our previous observation, we require $l(\mathcal{T}^*)$ to be at least $\frac{3(1+\epsilon')^2}{\epsilon'^4} = s$. So we add $\frac{s}{n}$ to each edge(both l_e and h_e) of the graph. This ensures that $l(\mathcal{T}^*)$ is at least s. Therefore, for the desired error guarantee ϵ, we have,

$$\begin{aligned}
&l(\mathcal{T}^{'}) + (s/n)n \geq (1-\epsilon)[l(\mathcal{T}^*) + (s/n)n] \\
&\Leftrightarrow l(\mathcal{T}^{'}) + s \geq (1-\epsilon)(l(\mathcal{T}^*) + s) \\
&\Leftrightarrow l(\mathcal{T}^{'}) \geq (1-\epsilon)l(\mathcal{T}^*) - s\epsilon \geq (1-\epsilon)l(\mathcal{T}^*) - \epsilon l(\mathcal{T}^*) = (1-2\epsilon)l(\mathcal{T}^*)
\end{aligned}$$

The last inequality follows from $l(\mathcal{T}^*) \geq s$. By choosing $\epsilon' = \epsilon/2$ we obtain the desired error bound, i.e., □

$$\mathbf{Pr}[l(\mathcal{T}^{'}) \leq (1-\epsilon')l(\mathcal{T}^*)] = c_1 < \frac{1}{e} \tag{3}$$

Using $\mu = \mathbb{E}[c(\mathcal{T}^{'})]$ and Equation 2

$$\mathbf{Pr}[c(\mathcal{T}^{'}) \geq B] \leq \mathbf{Pr}[c(\mathcal{T}^{'}) \geq (1+\epsilon')\mu] \leq \mathbf{exp}(-\frac{\epsilon'^2\mu}{3})$$

The second inequality follows from Chernoff bounds. Again if $\mathbf{exp}(-\frac{\epsilon'^2\mu}{3}) < \mathbf{exp}(-1)$, then $\mathbf{Pr}[c(\mathcal{T}^{'}) \geq B] < \frac{1}{e}$. Since $\mathbb{E}(c(\mathcal{T}^{'})) = \mu > \frac{3}{\epsilon'^2}$ can be achieved by scaling the edge costs c_e for each $h_e \in E$ along with the budget B. Note that l_e's continue to have 0 weights unlike the construction in Lemma 5. However, the basic argument can be extended by considering only the edge set of h_e.

$$\mathbf{Pr}[c(\mathcal{T}^{'}) \geq B] = c_2 < \frac{1}{e}. \tag{4}$$

The polynomial running time follow from the PTAS for bi-criterion spanning tree given by [RG96]. From Equations 3 and 4
$\mathbf{Pr}[l(\mathcal{T}^{'}) \geq (1-\epsilon)l(\mathcal{T}^*)\ ,\ c(\mathcal{T}^{'}) \leq B] \geq 1-c,$ where $(c = c_1 + c_2)$. This concludes the proof of Theorem 4. □

Corollary 6 *Given an undirected graph $G = (V, E)$ an error parameter ϵ the algorithm of section 2.2 returns a spanning tree $\mathcal{T}^{'}$ which is feasible within the budget B, such that*

$$\mathbf{Pr}[l(\mathcal{T}^{'}) \geq (1-\epsilon)l(\mathcal{T}^*)\ ,\ c(\mathcal{T}^{'}) \leq B] \geq 1 - \tfrac{1}{n},$$

in time polynomial in $|V| = n$, $|E| = m$ and $\frac{1}{\epsilon}$.

2.3 Further extensions

We now consider a more general version of the problem, where the edge lengths l_e instead of admitting only a single improvement h_e, can have multiple stages of improvement. That is for each $i \in [m]$, where m is the number of edges, we are given improved lengths $l_{i,1} \leq \cdots \leq l_{i,j(i)}$ with increasing improvement costs $c_{i,1} \leq \cdots \leq c_{i,j(i)}$. The lengths $l_{i,0}$ has cost $c_{i,0} = 0$ for each $i \in [m]$. The algorithm in section 2.2 can be extended to accommodate the multiple improvements in edge lengths. This is possible as each edge $i \in [m]$ in the multi-graph, that we constructed in section 2.2, can have $j(i)$ copies with respective costs $c_{i,1}, \cdots, c_{i,j(i)}$. Again, [RG96] would build a tree composing of two forests of improved and non-improved edge lengths. And then the improved

edge length forest can be sampled as in section 2.2 to ensure feasibility within the budget B. Approximation guarantees similar to the single stage improvement can be proved along similar lines.

Remark Our multistage improvement extend beyond the 0/1 upgradation model of [KNM$^+$97] but it doesn't directly yield results for the *integral* or the *rational* upgrade model of [KNM$^+$97]. However, we can approximate the rational model to within any inverse polynomial by choosing a polynomial number of parallel edges.

Due to the equivalence of maximum weight base and minimum cost base over a matroid, the results and analysis done in this section would also hold for the minimization version.

3 FPTAS for WILDAG

Given a DAG, $G = (V, E)$, having n vertices and m edges with edge lengths $l_i, i \in [m]$ where $[m] := \{1, ..., m\}$. Let $LDAG(l, s, t)$ denote the problem of finding the longest $s - t$ path in G, for a given source s and a sink t. An instance of the *Weight-Improvable-Longest-Path-in-DAG* is given by the same set of vertices $[n]$, and edges $[m]$ with edge-lengths $l_i, i \in [m]$, source s and sink t. Additionally, we are given the improved edge weights h_i for all $i \in [m]$, along with corresponding improvement costs q_i and an improvement budget B. A problem instance of *Weight-Improvable-Longest-Path-in-DAG* is denoted by $WILDAG(l, h, q, B, s, t)$. The optimum solution is denoted by $WILDAG^*(l, h, q, B, s, t)$.

Theorem 7 *The Weight Improvable Longest path problem in DAG is NP-hard.*

Proof: The Knapsack problem is also polynomial time reducible to the WILDAG problem using a construction similar to Figure 1. The two directed edges have associated tuples $(weight, cost)$ as $(0, 0)$ and (p_e, c_e) respectively where the knapsack problem is given by

$$\max \sum_e p_e x_e \quad s.t. \quad \sum_e c_e \leq B$$

Here B is the budget for improvement. □

3.1 A pseudo-polynomial algorithm for WILDAG

We can observe that if the query costs q_i for all $i \in [m]$ are uniform, say q, then the query costs and improvement budget B can be accordingly scaled, so that the improvement budget is basically the number of improvements allowed, say b. Also note that since G is a DAG, so the longest $s - t$ path can have at most $n - 1$ edges. Therefore $b \geq n - 1$ is as good as infinite budget, so the interesting case is only when $b < n - 1$.

We now give a dynamic programming formulation for $WILDAG$ with uniform query costs. Let $T =< v_1, ..., v_n >$ be a topological ordering of the vertices of the $DAG\ G$. Without loss of generality let us assume $v_1 = s$ and $v_n = t$. Let $L(v_i, q)$ denote the length of the longest $v_i - t$ path using the vertices $\{v_i, v_{i+1}, ..., t\}$ and at most q improved edge-weights.

$$L(v_i, q) = \max_{j:(i,j)\in E} \max\{L(v_j, q) + l_{e=(i,j)}, L(v_j, q-1) + h_{e=(i,j)}\} \tag{5}$$

This can be seen as follows. Let $S \subseteq T$ denote the set of vertices from which the longest path to t using at most q improved edges is known. The invariant for Equation 5 is that for each vertex $v \in S$, the length of the longest path from v to t using at most q improved edges is known. Thus the proof of correctness of 5 follows from the induction on size of S. Hence the longest path from s to t

using at most b improved edges is either by taking the improved edge h_e from s to v_i and taking the longest path from v_i to t using at most $b-1$ improved edges or by taking the un-improved edge l_e from s to v_i and taking the longest path from v_i to t using at most b improved edges, and maximizing over all neighbors v_i of s. Use the base cases as $L(t,q) = 0$ for all $q \in [b]$, and $L(v_i, x) = -\infty$ for all $i \in [m]$ and $x < 0$.

Lemma 8 *The dynamic programming algorithm takes $O(n^3)$ time.*

Proof: Each entry in the table can be computed in $O(n)$ time. The number of vertices in the DAG G is n and the number of improvements allowed is at most $n-2$. Hence the result follows. □ Now let us consider that the improvement costs q_e for all $e \in E$ that may not be uniform, and the improvement budget is B. We propose the following dynamic programming formulation. Let $L(v_i, w)$ denote the minimum improvement budget used such that the path from v_i to t using only the vertices $\{v_i, v_{i+1}...., t\}$ has length w. Let $W = \max_{e \in E} h_e$, then nW is a trivial upper bound on the length of the longest $s-t$ path in the DAG G. The dynamic programming recurrence follows along the lines of Equation 5

$$L(v_i, w) = \min_{j:(i,j)\in E} \min\{L(v_j, w - h_{e=(i,j)}) + q_e, L(v_j, w - l_{e=(i,j)})\} \tag{6}$$

Use the base cases as $L(t, w) = \infty$ for all $w > 0$, $L(v_i, x) = \infty$ for all $i \in [n]$ and $x < 0$, and $L(t, 0) = 0$.
The algorithm is presented below in Figure 2, and we prove the correctness in 3.1.1.

Algorithm `NonUni-Cost-Imp-DAG`(DAG G, $T =< v_1 = s, ..., v_n = t >$, B, $\bar{q}$,$\bar{l}$, $\bar{h}$)
Input: - DAG G,
- the topological ordering of the DAG G i.e, $T =< v_1, ..., v_n >$
- the improvement budget B
- the matrix of the improvement costs $\bar{q} = [q_1.....q_m]^T$
- the matrix of original edge weights $\bar{l} = [l_1.....l_m]^T$
- the matrix of improved edge weights $\bar{h} = [h_1.....h_m]^T$
Output: the length of the longest $s-t$ path spending at most budget B on improved edge-weights

Create the table $L = n \times nW$; $L(t, w) = \infty$ for all $w > 0$, $L(v_i, x) = \infty$ for all $i \in [n]$ and $x < 0$, and $L(t, 0) = 0$;
for $v_i = v_n, ..., v_1$ **do**
  for $w = 0, ..., nW$ **do**
  $L(v_i, w) =$ according to Equation 6
  end for
end for
return $\arg max_w\{L(s, w) \leq B\}$;

Figure 2: A pseudo polynomial time algorithm for WILDAG with non uniform query costs

3.1.1 Correctness and Running Time

The correctness of the algorithm in Figure 2 follows from the correctness of recurrence Equation 6. Let $S \subseteq T$ denote the set of vertices, such that the minimum improvement cost incurred on paths of length w from each vertex in S to t is known. Maintaining this invariant and inducting over the size of S gives the correctness of 6. The minimum improvement budget used such that the path

from s to t has length w is either incurring a cost q_e by taking the higher valued edge h_e to v_i and then incurring the minimum improvement budget such that the path from v_i to t (using vertices $\{v_i, \cdots, t\}$ in the topological ordering T) has length $w - h_e$, or not incurring any cost by taking the lower valued edge l_e to v_i and then incurring the minimum improvement budget such that the path from v_i to t has length $w - l_e$, and minimizing over all neighbors v_i of s. So, the length of the longest path form s to t is the largest w over all $s - t$ paths of length w such that the minimum improvement budget for the path is at most B.

Lemma 9 *The algorithm in Figure 2 takes $O(n^3W)$ time.*

Proof: Each entry in the table can be computed in $O(n)$ time. The number of vertices in the DAG G is n and the upper bound on the length of the longest $s - t$ path is nW. Hence the result follows. □ **Remark** For uniform cost, i.e., $q_e = 1$, we can rewrite the dynamic programming in terms of budget to obtain a polynomial time algorithm.

3.2 A FPTAS for WILDAG

Next, by scaling the lengths, we convert the previous pseudo polynomial time algorithm (in Figure 2) into an efficient version by compromising with an approximation factor in the objective of attaining the longest length $s - t$ path. Using $W = O(n/\epsilon)$, the running time of the algorithm in Figure 2 is $O(n^3.\frac{n}{\epsilon})$, which is similar to the classic FPTAS for Knapsack [Vaz01]. Thus the theorem follows.

Theorem 10 *The dynamic programming algorithm for WILDAG (Figure 2) with non uniform improvement costs returns a solution with objective value at least $(1-\epsilon)WILDAG^*$ in $O(\frac{n^4}{\epsilon})$ time.*

Remark In this section we have considered the problem of Weight Improvable Longest Path in a DAG but similar dynamic programming formulation extends to the problem of Weight Improvable Shortest Path in a DAG.

4 Concluding Remarks

It remains an open question if we can design a polynomial time algorithm for the uniform cost improvement version of the constrained MST. Further, it would be interesting to extend our techniques so that we can handle even continuous improvements on edge weights, more specifically, like the rational update model in [KNM+97].

Acknowledgement The authors wish to thank Marc Goerigk for valuable discussions and an anonymous reviewer for his suggestions on improving the presentation of an earlier version of this paper.

References

[AAN82] V. Aggarwal, Y.P. Aneja, and K.P.K. Nair. Minimal spanning tree subject to a side constraint. *Computers & Operations Research*, 9(4):287 – 296, 1982.

[BIO92] Oded Berman, DivinagraciaI. Ingco, and AmedeoR. Odoni. Improving the location of minisum facilities through network modification. *Annals of Operations Research*, 40(1):1–16, 1992.

[CLZ06] Ann Melissa Campbell, Timothy J. Lowe, and Li Zhang. Upgrading arcs to minimize the maximum travel time in a network. *Networks*, 47(2):72–80, 2006.

[DKNMR98] S.O. Drangmeister K.U., Krumke, H. Noltemeier, M.V. Marathe, and S.S. Ravi. Modifying edges of a network to obtain short subgraphs. *Theoretical Computer Science*, 203(1):91 – 121, 1998.

[DNW02] I. Demgensky, H. Noltemeier, and H.-C. Wirth. On the flow cost lowering problem. *European Journal of Operational Research*, 137(2):265–271, 2002. Graphs and Scheduling.

[FSO96] Greg N. Frederickson and Roberto Solis-Oba. Increasing the weight of minimum spanning trees. In *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '96, pages 539–546, Philadelphia, PA, USA, 1996. Society for Industrial and Applied Mathematics.

[GSSS14] Marc Goerigk, Yogish Sabharwal, Anita Schbel, and Sandeep Sen. Approximation algorithms for the weight-reducible knapsack problem. In T.V. Gopal, Manindra Agrawal, Angsheng Li, and S.Barry Cooper, editors, *Theory and Applications of Models of Computation*, volume 8402 of *Lecture Notes in Computer Science*, pages 203–215. Springer International Publishing, 2014.

[HN62] William W. Hardgrave and George L. Nemhauser. On the relation between the traveling-salesman and the longest-path problems. *Operations Research*, 10(5):pp. 647–657, 1962.

[KMR97] D. Karger, R. Motwani, and G.D.S. Ramkumar. On approximating the longest path in a graph. *Algorithmica*, 18(1):82–98, 1997.

[KNM+97] S.O. Krumke, H. Noltemeier, M.V. Marathe, S.S. Ravi, and K.U. Drangmeister. Modifying networks to obtain low cost trees. In Fabrizio d'Amore, PaoloGiulio Franciosa, and Alberto Marchetti-Spaccamela, editors, *Graph-Theoretic Concepts in Computer Science*, volume 1197 of *Lecture Notes in Computer Science*, pages 293–307. Springer Berlin Heidelberg, 1997.

[KNW+99] Sven O. Krumke, Hartmut Noltemeier, Hans-C. Wirth, Madhav V. Marathe, R. Ravi, S.S. Ravi, and R. Sundaram. Improving spanning trees by upgrading nodes. *Theoretical Computer Science*, 221(1–2):139–155, 1999.

[Kru56] J. B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. In *Proceedings of the American Mathematical Society, 7:48-50*, 1956.

[MRS+98] Madhav V. Marathe, R. Ravi, Ravi Sundaram, S. S. Ravi, Daniel J. Rosenkrantz, and Harry B. Hunt III. Bicriteria network design problems. *CoRR*, cs.CC/9809103, 1998.

[Phi93] Cynthia A. Phillips. The network inhibition problem. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Theory of Computing*, STOC '93, pages 776–785, New York, NY, USA, 1993. ACM.

[RG96] R. Ravi and M. X. Goemans. The constrained minimum spanning tree problem. In *Scandinavian Workshop on Algorithmic Theory (SWAT), Reykjavik, Iceland, Lecture Notes in Computer Science*, volume 1097, pp. 66-75, 1996.

[Vaz01] Vijay V. Vazirani. *Approximation Algorithms*. Springer-Verlag New York, Inc., New York, NY, USA, 2001.